

Exhibit 1

Price Discovery without Trading: Evidence from Limit Orders

JONATHAN BROGAARD, TERRENCE HENDERSHOTT, and RYAN RIORDAN*

ABSTRACT

We analyze the contribution to price discovery of market and limit orders by high-frequency traders (HFTs) and non-HFTs. While market orders have a larger individual price impact, limit orders are far more numerous. This results in price discovery occurring predominantly through limit orders. HFTs submit the bulk of limit orders and these limit orders provide most of the price discovery. Submissions of limit orders and their contribution to price discovery fall with volatility due to changes in HFTs' behavior. Consistent with adverse selection arising from faster reactions to public information, HFTs' informational advantage is partially explained by public information.

ACCORDING TO THE TRADITIONAL VIEW of price discovery, trades reveal investors' private information while market makers' quotes reflect public information (see, e.g., Glosten and Milgrom, 1985; Kyle, 1985). Most stock exchanges and financial markets have evolved into limit order books where there are no designated market makers and limit orders represent the bulk of activity. Theoretical models of limit order books study informed traders' choice between market orders and limit orders. The market/limit order choice of informed and uninformed investors determines the nature of price discovery and adverse selection. In this paper, we use regulatory data that enable the classification of limit orders and trades by high-frequency traders (HFTs) and non-HFTs to systematically quantify the contribution to price discovery of market and limit orders by HFTs and non-HFTs, primarily using a vector autoregression (VAR; Hasbrouck, 1991a, 1991b, 1995). We then link these results to theoretical models of limit order books.

*Jonathan Brogaard is with David Eccles School of Business, University of Utah. Terrence Hendershott is with Haas School of Business, University of California – Berkeley. Ryan Riordan is with Smith School of Business, Queen's University. The authors thank seminar participants at the 2015 Cambridge Microstructure Theory and Application Workshop, Australia National University, Baruch College, Boston College, Chinese University of Hong Kong, Goethe University, Hong Kong University, Stockholm Business School, UC Santa Cruz HFT Workshop, and University of Mannheim for helpful comments. The authors also thank Helen and Victoria asked not to be thanked IIROC for providing data and comments. All errors are our own. This research was supported by the Social Sciences and Humanities Research Council of Canada and the Norwegian Finance Initiative. Hendershott has provided expert witness testimony in a variety of matters, including an ongoing market manipulation case.

DOI: 10.1111/jofi.12769

The role of HFTs in adversely selecting non-HFTs is widely debated by academics, regulators, and investors, with concerns often focusing on HFTs using market orders to “pick off” stale limit orders (Biais, Foucault, and Moinas, 2015; Budish, Cramton, and Shim, 2015; Foucault, Hombert, and Rosu, 2015; Foucault, Kozhan, and Tham, 2017). Relatedly, most empirical literature on price discovery focuses on the contribution of market orders (Hasbrouck, 1991a, 1991b; Brogaard, Hendershott, and Riordan, 2014).¹ We find, however, that HFTs’ limit orders contribute more than twice as much to price discovery as their market orders. In contrast, non-HFTs’ market orders contribute more to price discovery than their limit orders. Comparing HFTs to non-HFTs, we find that HFTs’ market orders are responsible for less price discovery than non-HFTs’ market orders, while HFTs’ limit orders are responsible for twice as much price discovery as non-HFTs’ limit orders. Overall, HFTs’ market orders play a smaller role in price discovery while HFTs’ limit orders play a larger role.²

Our results show that more aggressive orders have a higher price impact.³ Market orders, the most aggressive order type, have the highest impact, followed by orders that change the national best bid and offer (NBBO), orders at the NBBO, and orders behind the NBBO. Despite their lower individual price impact, limit orders provide the majority of price discovery because they are far more numerous: market orders represent less than 5% of messages.⁴ Individual HFT market orders contribute more to price discovery on average than

¹ Recent empirical literature examines the contribution of both market orders and limit orders to price discovery. Among others, Hautsch and Huang (2012) quantify the impact of a limit order using a cointegrated VAR. Cao, Hansch, and Wang (2009) and Cont, Kukanov, and Stoikov (2014) show that order imbalances predict future price movements. Fleming, Mizrach, and Nguyen (2017) study the price impact of market and limit orders in U.S. Treasury bonds using a VAR setup similar to the one used here. These papers focus on the price impact of individual orders. In contrast, we focus on decomposing the total amount of price discovery. In addition, our data identifying HFTs allow us to examine whether orders placed by different traders play different roles in price discovery.

² These results do not necessarily imply that concerns about HFTs are unwarranted. HFTs’ market orders contribute some to price discovery and adverse selection. HFTs incorporating information with limit orders can also cause non-HFTs’ limit orders to be adversely selected, which could in turn lead to excess intermediation if non-HFTs reduce their use of limit orders (Jovanovic and Menkveld, 2015). However, the relative magnitudes of HFTs’ limit orders and market orders in contributing to price discovery suggests that HFTs are not primarily using their information in market orders to adversely select non-HFTs.

³ A number of papers examine HFTs’ trading and price discovery. Brogaard, Hendershott, and Riordan (2014) show that HFTs contribute to price discovery with market orders by trading in the direction of future price changes. Carrion (2013) finds that market-wide information is incorporated into prices quickly on days when HFTs trade more. Conrad, Wahal, and Xiang (2015) find that high-frequency trading and quoting correlate with more efficient prices. Chaboud et al. (2014) find that HFTs improve price efficiency through lower return autocorrelations and fewer arbitrage opportunities. Chordia, Green, and Kottimukkalur (2018) show that high-frequency market orders impound information quickly following macroeconomic announcements.

⁴ A *message* refers to any instruction received by an exchange; a message includes marketable orders, limit order placements, limit order cancellations, and limit order replacements. A *limit order* refers to a nonmarketable instruction received by an exchange; a limit order includes limit order placements, limit order cancellations, and limit order replacements.

non-HFT market orders. However, non-HFT market orders are three times more likely, making them more important overall. HFT limit order submissions and cancellations are both frequent and informative, leading HFT limit orders to contribute roughly 30% of total price discovery versus roughly 15% for non-HFTs.

That HFT limit order submissions have a positive price impact is seemingly at odds with the results in Brogaard, Hendershott, and Riordan (2014) that suggest HFTs' liquidity-supplying trades have a negative price impact.⁵ However, the analysis in Brogaard, Hendershott, and Riordan (2014) relies on executed trades and hence does not capture the effect of limit orders that do not execute. For example, a limit order to buy will not execute if the price increases. In this case the limit order contributes to price discovery without trading. When a limit order executes, in contrast, its price impact is effectively the opposite of the market order that it executes against. For example, when a buy market order executes against a sell limit order, on average the efficient price will increase. This leads to the buy market order having a positive price impact and the sell limit order having a negative price impact upon execution. This is why Brogaard, Hendershott, and Riordan (2014) find that HFTs' liquidity-supplying trades have a negative price impact. The price impact of limit orders upon submission is the weighted average of the (negative) price impact of the limit orders that execute and the (positive) price impact of the limit orders that do not execute. Given that only 5% of limit orders execute, it is not surprising that the average price impact for executed and nonexecuted limit orders is positive.

Theoretical models of limit order books provide insights into the roles that different orders by different traders play in price discovery (e.g., Goettler, Parlour, and Rajan, 2009 [GPR]; Hoffmann, 2014).⁶ These models focus on traders' choice between market orders and limit orders based on traders' information and the state of the limit order book. Limit orders receive rather than pay the bid-ask spread, but do not execute with certainty. Market orders always execute, but pay the bid-ask spread. When information is more valuable and the spread is narrower, traders prefer market orders to limit orders.

⁵ Brogaard, Hendershott, and Riordan (2014) show that HFTs contribute to price discovery with market orders by trading in the direction of future price changes. They also find that HFTs' liquidity-supplying trades are in the opposite direction of future price changes. Their results are not necessarily inconsistent with the results presented here. We show that aggressive limit order submissions and cancellations are associated with positive price impacts at the time of submission. Brogaard, Hendershott, and Riordan (2014) show that orders submitted by HFTs that are not subsequently cancelled and that execute against more aggressively priced incoming limit orders are adversely selected.

⁶ Other papers examine limit order trading by informed investors, but provide insights less directly related to HFTs. Kaniel and Liu (2006) theoretically model the order choice of informed traders. Consistent with GPR, their two-period model finds that informed traders are more likely to submit limit orders. Bloomfield, O'Hara, and Saar (2005) conduct an experiment that includes long-lived private information and show that informed trades submit more limit orders. Similarly, Rosu (2019) shows that informed traders tend to use limit orders for moderate levels of mispricing and market orders more extreme mispricing.

Traders with different characteristics face different trade-offs in their order choice. Traders with no intrinsic motivation to trade (GPR) and fast traders (Hoffmann, 2014) who can revise their orders more often prefer limit orders because execution uncertainty is less costly for them and their limit orders face less adverse selection, respectively. HFTs fit both of these descriptions. These models are consistent with our empirical finding that HFTs submit the majority of limit orders.⁷ Also consistent with our empirical results, GPR and Hoffmann (2014) find that limit orders play a significant role in price discovery.⁸

GPR and Hoffmann (2014) study closely related models in which traders choose between market and limit orders and later-arriving traders observe new public information.⁹ While GPR and Hoffmann (2014) find similar results for limit orders overall and for traders with characteristics shared by HFTs, their models yield different predictions for how their results vary with volatility.¹⁰ An important difference between GPR and Hoffmann (2014) is that in GPR, investors have different and often large private gains from trade.

The basic trade-off in both models is between the risk of nonexecution and the adverse-selection risk associated with the submission of limit orders. When volatility is high, the picking-off risk is higher. Hoffmann's (2014) fast traders can avoid being adversely selected by slow traders but cannot submit market orders profitably. Therefore, fast traders increase their limit order submissions when volatility is high.¹¹ GPR include investors with large private values and speculators with zero private value. When volatility is high, extreme private value investors submit better-priced limit orders to entice speculators to submit market orders.¹² The speculators in GPR then switch from limit to market orders. We find that HFTs reduce their use of limit orders when volatility

⁷ Jovanovic and Menkveld (2015) model a continuously present intermediary that can constantly revise its limit orders as "public" information arrives. This intermediary (HFT) also submits more limit orders than market orders.

⁸ Prior papers empirically examine limit order usage by non-HFT traders. Collin-Dufresne and Fos (2015) show that one group of informed traders, namely 13D activist investors, use limit orders. Anand, Chakravarty, and Martell (2005) show that institutions use limit orders. Our results on non-HFT limit orders contributing to price discovery are consistent with the use of limit orders by traders with long-lived information. Using the same regulatory data as in our paper, Korajczyk and Murphy (2019) provide some evidence of informed institutions using limit orders.

⁹ Foucault (1999) provides many of the basic building blocks in GPR and Hoffmann (2014). For example, in these models adverse selection arises from the arrival of public information and limits orders not always being immediately cancelable. However, all traders are the same in Foucault (1999), which provides limited insight into HFTs and price discovery.

¹⁰ Few models examine fast, informed traders' use of limit or market orders conditional on volatility. An exception is Baldauf and Mollner (2016). Similar to Hoffmann (2014), their setting allows traders to increase their trading speed, for a fee. They model a dynamic setting with information arrivals. Their model suggests that fast, informed traders are more likely to use market orders when the degree of mispricing is high and limit orders when the degree of mispricing is low.

¹¹ Relative to fast traders, slow traders in Hoffmann (2014) submit relatively more limit orders than market orders when volatility is high. This is consistent with the model of Foucault (1999), in which investors are homogeneous and more limit orders are submitted relative to market orders when volatility is high.

¹² Bloomfield, O'Hara, and Saar (2005) also study volatility and limit orders. They analyze two types of variability: volatility and extremity. Volatility is captured by the distribution of future

is high.¹³ This suggests that modeling richer heterogeneity (large, small, and zero) in the private valuations of trading motives, as in GPR, is important for understanding HFTs and limit-order markets.

Price discovery switching from limit orders to market orders has implications for market stability and possible market failure.¹⁴ Madhavan (1992) shows that continuous markets can fail when adverse selection is sufficiently high. If volatility increases due to greater private information, then market failure is more likely. If informed liquidity providers switch from limit orders to market orders when volatility increases, then market failure is even more likely. We show that HFTs reduce their use of limit orders as volatility increases, which reduces the contribution of their limit orders to price discovery. When volatility increases, the market/limit order trade-off between execution speed/certainty and price increases more in favor of market orders for informed traders than uninformed traders. This differential trade-off for informed and uninformed traders raises concerns that endogenous fragility in continuous limit order books. This question represents an important area for future empirical and theoretical research.

Prior literature raises several concerns about HFTs. First, a number of theoretical papers show that fast traders like HFTs can adversely select slower traders. For example, Foucault, Hombert, and Rosu (2015), Biais, Foucault, and Moinas (2015), and Budish, Cramton, and Shim (2015) show that some traders trading faster on public signals increases information asymmetry.¹⁵ Our results lend some support to these concerns. For example, we find that the larger price impact of HFTs' orders is explained in part by public information, such as the state of the limit order book and lagged returns in a correlated asset (e.g., the TSX 60 exchange traded fund). However, our results also suggest that trading on such public information is not the dominant role of HFTs in overall price discovery.¹⁶ Second, extant work suggests that HFTs could be

values whereas extremity is captured by the distance between the previous price and the future value. For distributional volatility conditional on extremity, Bloomfield et al. (2005) find no relationship between limit order submissions and volatility. For their extremity measure of volatility, Bloomfield et al. (2005) find a negative relationship between volatility and the limit order submissions of informed trades. This suggests that long-lived private information may also be able to generate the negative empirical correlation between volatility and limit order.

¹³ Ahn, Bae, and Chan (2001) show that limit order submissions increase subsequent to increases in transitory volatility.

¹⁴ Danielsson, Shin, and Zigrand (2012) and Kirilenko et al. (2017) discuss endogenous extreme events and the 2010 flash crash. Brogaard, Hendershott, and Riordan (2018) show that the trades of HFTs supply liquidity more than they demand liquidity during extreme price movements. Unlike our examination of how HFTs' limits orders change with volatility, Brogaard, Hendershott, Riordan (2018) compare levels during extreme price movements without controlling for the relative liquidity supplied and demanded by HFTs outside of extreme price movements. If HFTs supply liquidity more than they demand liquidity, which is true in our sample, HFTs could decrease their use of limit orders as volatility increases while their trades still supply liquidity more than demand liquidity.

¹⁵ For empirical evidence, see Brogaard, Hendershott, and Riordan (2018).

¹⁶ Brogaard, Hendershott, and Riordan (2014) find that HFTs' trading correlates with public information in past market-wide stock returns and limit order book imbalances. We find that this is

“front-running” non-HFTs’ orders by detecting large non-HFT orders that are split over time or across exchanges (Hirschey, 2017; Korajczyk and Murphy, 2019; van Kervel and Menkveld, 2019). We find that HFTs’ orders generally do not anticipate non-HFTs’ orders in the same direction and that HFTs may react more quickly to public information. For instance, HFT orders that move the NBBO negatively predict the same non-HFT orders and non-HFT market orders at the NBBO. While this finding is consistent with HFTs observing and reacting to public signals before non-HFTs are able to react, we also show that the contribution of HFTs’ limit orders to price discovery is not due solely to their orders arriving only slightly ahead of non-HFTs’ limit orders.

We also examine price discovery for stocks across markets. In particular, we examine HFTs’ activity and price discovery on each exchange and across exchanges. As with the market-wide results, we find that HFTs are the predominant channel of price discovery on each exchange through their limit orders. While significant price discovery occurs within the same second across exchanges, the role of HFTs’ limit orders for price discovery is predominant even when sampling at the one-second frequency and this does not appear to be due solely to minuscule differences in speed. Finally, HFTs react more to events on other exchanges than non-HFTs, this is consistent with HFTs integrating information across markets. While this seems beneficial in a fragmented market, whether such integration is better than trading on a centralized exchange is an open question.

The remainder of the paper proceeds as follows. Section I describes the data and institutional details. Section II documents market-wide activity and price discovery. Section III provides evidence on within- and across-market activity and price discovery. Section IV concludes.

I. Data and Institutional Details

Data are provided by the Investment Industry Regulatory Organization of Canada (IIROC). The data include every message submitted on recognized equity markets in Canada with masked market IDs, masked participant IDs, security IDs, date and timestamps to the millisecond, order type, order volume, and a buy/sell indicator.¹⁷ Importantly, the data identify activities across exchanges as the masked participant IDs remain constant across days, securities,

true for HFTs’ limit orders submissions as well. Whether to attribute any HFT contribution to price discovery to private information runs into the deeper issue described in Hasbrouck (1991a, 190): “the distinction between public and private information is more clearly visible in formal models than in practice.” Given that HFTs rely only on public information in their trading algorithms, one can argue that all of their contribution to price discovery is due to public information. However, if HFTs’ algorithms better interpret public signals (like the short sellers on news days in Engelberg, Reed, and Ringgenberg, 2012), then it is more difficult to characterize HFTs as incorporating purely public information.

¹⁷ The data are structured similar to the NASDAQ ITCH. They contain every message sent by each participant to the exchange. The messages include the initial order, cancels, and amendments to the order. As in the United States, there are a number of different order types, such as hidden orders and immediate or cancel (IOC) orders, which are flagged in the data. We exclude hidden limit

and markets. IIROC requires exchanges to report messages in a standardized format. As such, some order types may be recorded as multiple orders. For example, modifications are reported to IIROC as both a cancel and a new order.

A. Trading Landscape

Canada has a number of equity markets on which trading is organized. We identify nine in total and present summary statistics on the three largest exchanges.¹⁸ Trading on these three exchanges makes up more than 98% of the total trading volume in our sample stocks over our sample period.

Markets in Canada are similar to U.S. markets in that electronic limit order books observe price-time-display priority. Canadian markets differ from U.S. markets during our sample period in that Canadian markets are less fragmented, do not allow subpenny trading, and regulate that dark orders improve the price by half of a tick. Orders in Canada during the sample period are protected via order protection rules (OPR).¹⁹ OPR apply to marketplaces that provide “automated functionality.” Automated functionality includes automatically displaying and updating the status of each participant’s orders, as well as immediately and automatically accepting incoming orders, executing those orders, and canceling any unexecuted portion of those orders marked as immediate-or-cancel (IOC). OPR apply only to visible orders and the visible parts of orders and require marketplaces to implement rules to prevent trade-throughs, that is executing before “immediately accessible, visible, better-priced limit orders.”

In contrast to Regulation NMS in the United States, Canadian markets implement full depth-of-book protection. This means that before an order is executed, marketplaces must ensure that all protected orders that are visible at better price levels have been executed. Canadian regulations also impose best execution obligations on brokers. These regulations require dealers and advisors “to execute a trade on the most advantageous terms reasonably available under the circumstances when acting for a client.” See Korajczyk and Murphy (2019) for additional institutional details.

The Canadian market has seen a dramatic increase in competition for investor order-flow since 2008. In May of 2007 a consortium of Canada’s largest banks announced a trading platform designed to compete with the TSX, called Alpha Trading Systems. Shortly thereafter in December of 2007 Chi-X

orders from the order book construction. IOC orders are included as an order and cancel if they are not executed, and a trade if they are filled. IIROC receive data with homogenized fields from each exchange in a format that allows for cross-platform integration. Specifically, exchange data must follow the Financial Information Exchange (FIX) protocol (<http://www.fixtradingcommunity.org/>). Any deviation from the FIX implementation must be approved by IIROC with a regulatory-feed-compliant solution. The data are timestamped by each exchange. The exchanges are required to synchronize their clocks with IIROC, which follows the National Research Council Cesium Clock.

¹⁸ For an overview of marketplaces as of June 1, 2015, see http://www.iroc.ca/industry/marketmonitoringanalysis/Documents/SumCompEquityMarkets_en.pdf.

¹⁹ See http://www.osc.gov.on.ca/en/Marketplaces_order-protection_index.htm.

announced their intention to commence trading in selected Canadian stocks on February 20, 2008. In response TSX rolled out new trading technology (TSX Quantum) to all TSX-listed stocks. In 2012 TSX's parent company, the Maple Group, purchased Alpha and now operates Alpha as a separate exchange within the TMX group of exchanges.

B. Sample

Our sample comprises the 15 securities that are part of the TSX 60, the primary Canadian equity index, at the end of 2014 that are not cross-listed in the United States; the other 45 stocks in the TSX 60 are cross-listed. We exclude cross-listed stocks as we cannot measure message activity that occurs off Canadian exchanges in the same way. In addition, cross-listed stocks may have different properties (Bacidore and Sofianos, 2002). Table I reports descriptive statistics for the sample stocks: market capitalization, share price, trade size, number of trades, number of shares traded, dollar volume traded, NBBO quoted half-spread, % HFT, % HFT demand, % HFT supply, and the standard deviation of returns. The average market capitalization from October 15, 2012 to June 28, 2013, the sample period, is *Market Cap*,²⁰ while the daily standard deviation of stock returns based on end-of-day prices during the sample period is *Std. Dev. of Returns*. Market capitalization and the standard deviation of stock returns are based on data from Datastream. All other variables are reported as stock-day averages during the sample period using IIROC data. Table I includes activities from all exchanges, whereas the remaining tables include observations only from the three largest exchanges.

The firms in our sample have market capitalization that ranges from \$1.95 billion CAD to more than \$28 billion CAD. Share prices vary between \$20 and \$76, with the exception of Bombardier, with a price of \$4.00. The stocks in our sample are actively traded with between \$11.84 million and \$70.97 million traded per stock-day. The stocks are relatively liquid with quoted half-spreads between 1.38 and 12.64 basis points.

C. HFT Classification

We classify trader IDs as HFTs using the following criteria over the entire sample period:

- (i) make up more than 0.25% of trading volume;
 - (ii) have an end-of-day inventory of less than 20% of their trading volume;
- and

²⁰ The sample starting date is just after Canadian regulators began requiring dark liquidity provision to improve on the best displayed prices by at least 1 cent, or 1/2 cent if the displayed spread is 1 cent. To examine whether our results are sensitive to slow adjustment to this regulatory change, Internet Appendix Tables AII to AVI repeat the main analyses for the 2013 subsample and show economically similar results. The Internet Appendix is available in the online version of the article on the *Journal of Finance* website.

Table I
Descriptive Statistics

The table reports summary statistics for the 15 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. *Ticker* is the ticker. *Market Cap.* is the average market capitalization from Datastream, in billions of Canadian dollars. *Share Price* is the average traded stock price. *Trade Size* is the average trade size in dollars. *Number of Trades* is the average number of trades, in thousands. *Number of Shares Traded* is the average number of shares traded, in thousands. *Dollar Volume Traded* is the number of shares traded multiplied by the stock price, in millions of dollars. *NBBO Quoted Half-Spread* is the calendar time weighted one-half quoted difference between the national best bid and the national best ask price, in basis points. *% HFT* is the double-sided dollar volume percentage of trades by a high-frequency trader (HFT). *% HFT Demand* is the dollar volume percentage of trades in which an HFT is the liquidity taker. *% HFT Supply* is the dollar volume percentage of trades in which an HFT is the liquidity provider. *Std. Dev. of Returns* is the standard deviation of the daily returns for the stock, in percent.

Ticker	Market Cap. (\$ Billion)	Share Price (2)	Trade Size (3)	Number of Trades ('000) (4)	Number of Shares Traded ('000) (5)	Dollar Volume Traded (\$ Million) (6)	NBBO Quoted Half-Spread (bps) (7)	% HFT (8)	% HFT Demand (9)	% HFT Supply (10)	Std. Dev. of Returns (%) (11)
ARX	\$8.00	\$25.59	\$4,854	5.47	1,053.84	\$26.91	2.72	21.6%	15.0%	28.1%	1.23
ATD.B	\$10.03	\$53.40	\$10,517	2.28	446.81	\$24.03	4.07	17.4%	22.9%	12.0%	1.17
BBD.B	\$6.95	\$4.00	\$4,875	10.14	12,697.71	\$50.42	12.64	30.1%	19.5%	40.6%	2.00
COS	\$9.90	\$20.43	\$4,765	9.13	2,129.80	\$43.52	2.52	24.9%	14.4%	35.4%	1.27
CTC.A	\$6.77	\$72.35	\$14,821	1.67	343.34	\$25.12	4.37	15.3%	19.1%	11.6%	1.51
FM	\$10.95	\$19.77	\$5,383	13.41	3,656.19	\$70.97	2.92	18.4%	14.2%	22.7%	2.71
FTS	\$6.96	\$33.63	\$6,943	3.91	791.23	\$26.52	1.84	23.5%	13.9%	33.2%	0.77
HSE	\$28.74	\$29.24	\$6,307	7.28	1,545.42	\$45.27	2.18	24.4%	19.0%	29.8%	1.28
L	\$11.62	\$41.10	\$9,738	2.99	783.63	\$33.09	2.89	12.1%	11.4%	12.8%	1.56
MRU	\$1.95	\$52.49	\$10,175	2.11	3.97	\$20.87	4.11	16.1%	16.6%	15.6%	0.85
NA	\$12.46	\$76.46	\$14,217	4.31	797.08	\$60.71	1.38	21.4%	18.7%	24.1%	0.62
POW	\$12.10	\$26.28	\$6,210	5.25	1,224.45	\$32.39	2.17	23.0%	14.4%	31.6%	0.92
SAP	\$9.61	\$48.92	\$11,088	1.99	421.56	\$20.79	3.81	20.3%	23.0%	17.7%	0.97
SNC	\$6.45	\$42.50	\$10,066	3.22	756.68	\$32.31	3.70	14.4%	16.4%	12.3%	1.54
WN	\$9.35	\$72.98	\$12,766	0.91	160.60	\$11.84	5.77	15.8%	22.2%	9.5%	1.08
Average	\$10.12	\$41.28	\$11,637	4.94	1,787.49	\$34.98	3.81	19.9%	17.4%	22.5%	1.30

- (iii) never hold more than 30% of their daily trading volume at one time within the trading day.

The methodology uses similar trade and inventory characteristics as Kirilenko et al. (2017), who also categorize traders as HFTs based on data-driven criteria. While there is no commonly accepted definition for all characteristics of HFTs, these criteria capture an important aspect of HFTs: they are relatively large short-term speculators. Small HFTs or HFTs that hold larger positions would be classified as non-HFTs. Other strategies that could be thought of as corresponding to HFTs, for example, arbitraging between the cash and futures market where offsetting positions are often held overnight, could also be incorrectly characterized as corresponding to non-HFTs. In addition, trading firms can have multiple IDs. If HFT firms split their trading across IDs, such that the individual IDs fail to meet our criteria but the aggregation across all firm IDs would meet the criteria, then our approach would incorrectly identify those IDs as non-HFT. Any misclassifications would attenuate the differences between HFTs and non-HFTs. We identify 61 HFT IDs from the 1,706 IDs in the Canadian market. Each stock-day an average of 13.06 HFT IDs are active while 185.36 non-HFTs are active. When HFTs are active in a stock, they trade on average on 2.06 exchanges per stock-day.

Table I shows that HFT participation varies across the sample of stocks, ranging from 12.1% to 30.1%. HFT liquidity demand and supply participation in trades are not evenly distributed, ranging from 11.4% and 23.0% for demand and 9.5% and 40.6% for supply. HFTs generally supply more liquidity than they demand in our sample, but there is variation across stocks, suggesting that HFT strategies may also vary across stocks.

Table II reports statistics on HFT and non-HFT participants. The average HFT is more active in terms of message activity, trades, and shares and volume traded, and has a higher message-to-trade ratio.²¹ Overall, HFTs hold less inventory throughout and at the end of the trading day. HFTs hold considerably less inventory than their trading would imply.

HFTs on average submit 15 times as many messages (4,450 vs. 290) and trade six times more often (240 vs. 41) than non-HFTs. HFTs' order-to-trade ratio is more than twice that of non-HFTs (65 vs. 27), and HFTs' average trade size is less than one-sixth the size of non-HFTs' trades (\$5,664 vs. \$37,532). Overall, we find that HFTs exhibit small trade sizes, high order-to-trade ratios, and other commonly accepted characteristics of HFTs, which provides support for our classification approach.

We also compile inventory statistics to better understand how HFTs and non-HFTs manage their inventory.²² We report the absolute value of the end-of-day inventory/dollar volume traded for the given stock-day, the absolute value of the maximum inventory observed/dollar volume traded for the stock-day, and the number of days on which the absolute value of the end-of-day

²¹ These values may be sensitive to outliers. See Internet Appendix Table AI for Table II descriptive statistics based on median values rather than means.

²² The inventory measure assumes that each ID begins each day with a zero position.

Table II
Trader Type Statistics

The table reports stock-day-participant (HFT and non-HFT) average trading, orders, and positions for individual HFT and non-HFTs for the 15 non-cross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. *Number of Messages* is the number of orders, order cancels, and order amends that a trader places. *Number of Trades* is the number of trades (marketable orders) conducted by a trader. *Number of Shares Traded* is the number of shares traded by a trader. *Dollar Volume (DV) Traded* is the number of shares traded by a trader multiplied by the share price. *Message-to-Trade Ratio* is the number of messages deployed for each trade by a trader. *DV Traded/Total DV Traded* is the dollar volume traded by a trader scaled by the total dollar volume traded on that stock-day. *Abs(EoD Inv.)/DV Traded* is the absolute value of a trader's end-of-day dollar volume inventory scaled by that trader's dollar volume traded. *Abs(Max Intra. Inv.)/DV Traded* is the trader's absolute value of the maximum intraday dollar volume inventory position scaled by that trader's dollar volume traded. *% of days with Abs(EoD Inv.)/DV Traded < 3%* is the percentage of stock-day-trader observations with *Abs(EoD Inv.)/DV Traded* less than 3%. *Average DV Trade Size* is the average dollar volume size of a trade. *Number of Participants* is the average number of traders in each stock-day.

	HFT (1)	Non-HFT (2)
Number of Messages (thousand)	4.45	0.29
Number of Trades	239.58	41.05
Number of Shares Traded (thousand)	68.96	17.44
Dollar Volume (DV) Traded (\$million)	\$1.03	\$0.35
Message-to-Trade Ratio	64.74	27.01
Number of Messages/Total Messages	2.81%	0.21%
DV Traded/Total DV Traded	1.51%	0.46%
Abs(EoD Inv.)/DV Traded	10.68%	69.82%
Abs(Max Intra. Inv.)/DV Traded	17.99%	79.63%
% of days with Abs(EoD Inv.)/DV Traded < 3%	50.21%	6.64%
Average DV Trade Size	\$5,663.56	\$37,531.60
Number of Participants	13.06	185.36

inventory/dollar volume traded is below 3%. For all inventory statistics, HFTs hold less inventory relative to their trading volume at the end of the day, 11% versus 70%, and have lower intraday maximums, 18% versus 80%. HFTs' inventories are consistent with short-run speculators closely managing risk.

To begin our analysis of the overall activities of HFTs and non-HFTs, in Table III we report the frequency of messages by participant type, order type, and aggressiveness. Limit order aggressiveness is determined relative to the NBBO: marketable limit orders (trades) that change the NBBO and those that do not, limit orders and cancels that move the NBBO, limit orders and cancels at the NBBO, limit orders and cancels one tick behind the NBBO, and limit orders and cancels more than one tick behind the NBBO. On the average stock-day, 72,466 orders, including trades and cancelations, are placed. Table II shows that individual HFTs are much more active, but the large number of non-HFTs results in HFTs comprising roughly half (53%) of aggregate message activity.²³

²³ Tables I and II correspond to the full sample. See Table III and thereafter the sample that meets our requirements for the VAR in Tables IV and V.

Table III
Message Frequency

The table reports the frequency of orders by type-aggressiveness broken down by participant for the 15 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. This analysis considers all messages on the three exchanges by HFTs and non-HFTs that are Trades, Orders, or Order Cancels. *Trade - Change Price* captures trades that consume liquidity beyond the NBBO. *Trade - Same Price* captures trades that do not consume liquidity beyond the NBBO. *Improving Order* captures orders that move the NBBO. *Order Placement at NBBO* captures orders at the NBBO. *Order 1 tick from NBBO* captures orders one cent away from the NBBO. *Order > 1 tick from NBBO* captures orders more than one cent away from the NBBO. For cancels, the analogous definitions apply. *Total Number of Observations* is the number of observations on all three exchanges on the average stock-day.

	HFT (1)	Non-HFT (2)
Trade - Change Price	0.21%	0.28%
Trade - Same Price	0.76%	3.40%
Improving Order	2.15%	1.29%
Worsening Cancel	0.89%	0.63%
Order Placement at NBBO	14.88%	8.45%
Order Cancel at NBBO	12.42%	5.44%
Order 1 tick from NBBO	3.85%	1.55%
Cancel 1 tick from NBBO	4.61%	2.56%
Order > 1 tick from NBBO	6.33%	9.88%
Cancel > 1 tick from NBBO	7.33%	13.09%
Total Number of Observations	72,466	

Limit order submissions are the most frequent messages, making up 48%, with cancellations making up 47%. Market orders, the focus of most of the literature on price discovery, make up less than 5% of all activity. Activity at and better than the NBBO makes up roughly 50% of overall activity. HFTs are more active at best prices, making up 60% of total activity, and less active than non-HFTs at prices at least one tick from the NBBO, 22% versus 27%. Activity that moves the NBBO is likely important for price discovery, but makes up only 5% of all market activity: 0.5% market orders, 3.4% limit order submissions, and 1.5% limit order cancellations. Order from HFTs that move the NBBO are almost twice as frequent as the same orders from non-HFTs, 2.15% versus 1.29%. The remaining 95% of activity does not change prices. Below we show that these orders do not play an important role in price discovery.

II. Market-Wide Price Discovery

To investigate how different messages impact price discovery, we use Fleming, Mizrach, and Nguyen’s (2017) extension of Hasbrouck (1991a). The VAR model estimates movements in the efficient price using past price movements and the arrivals and cancellations of new orders. More specifically, the model dynamically estimates the relation between order submissions and

cancellations and future movements in the efficient price.²⁴ Causal interpretations in the VAR come from the ordering of events in time and do not account for possibly small time differences in reporting across exchanges. In Section III we examine calendar-time measures of price discovery that are less sensitive to such small time differences. The VAR incorporates trades and limit order activity, both submissions and cancellations, at various price levels into the standard price discovery VAR. We extend this model further by separating HFTs' and non-HFTs' activity. For each stock-day, the VAR model that we estimate is:

$$\begin{aligned}
 r_t &= \sum_{i=1}^5 \alpha_i^1 r_{t-i} + \sum_{i=0}^5 \beta_i^{1,1} X_{t-i}^1 + \sum_{i=0}^5 \beta_i^{1,2} X_{t-i}^2 + \cdots + \sum_{i=0}^5 \beta_i^{1,20} X_{t-i}^{20} + \mu_t^1 \\
 X_t^1 &= \sum_{i=1}^5 \alpha_i^2 r_{t-i} + \sum_{i=1}^5 \beta_i^{2,1} X_{t-i}^1 + \sum_{i=1}^5 \beta_i^{2,2} X_{t-i}^2 + \cdots + \sum_{i=1}^5 \beta_i^{2,20} X_{t-i}^{20} + \mu_t^2 \\
 X_t^2 &= \sum_{i=1}^5 \alpha_i^3 r_{t-i} + \sum_{i=1}^5 \beta_i^{3,1} X_{t-i}^1 + \sum_{i=1}^5 \beta_i^{3,2} X_{t-i}^2 + \cdots + \sum_{i=1}^5 \beta_i^{3,20} X_{t-i}^{20} + \mu_t^3 \\
 &\vdots = \vdots \\
 X_t^{20} &= \sum_{i=1}^5 \alpha_i^{21} r_{t-i} + \sum_{i=1}^5 \beta_i^{21,1} X_{t-i}^1 + \sum_{i=1}^5 \beta_i^{21,2} X_{t-i}^2 + \cdots + \sum_{i=1}^5 \beta_i^{21,20} X_{t-i}^{20} + \mu_t^{21},
 \end{aligned}$$

where α is the coefficient on the midpoint return series, r , lagged one to five periods and β is the coefficient on the 20 limit order and trade variables, $X^1 - X^{20}$. Note that contemporaneous events can cause price changes: the return equation includes contemporaneous and five lag values of the limit order and trade variables, whereas the remaining VAR equations do not include contemporaneous terms.

The VAR is in event time, t , with each message being an observation. The 20 X variables are as follows: *HFT Trade – Change Price*, *HFT Trade – Same Price*, *HFT Improving Order*, *HFT Worsening Cancel*, *HFT Order Placement at NBBO*, *HFT Cancel at NBBO*, *HFT Order 1 tick from NBBO*, *HFT Cancel 1 tick from NBBO*, *HFT Order > 1 tick from NBBO*, *HFT Cancel > 1 tick from NBBO*, *non-HFT Trade – Change Price*, *non-HFT Trade – Same Price*, *non-HFT Improving Order*, *non-HFT Worsening Cancel*, *non-HFT Order Placement at NBBO*, *non-HFT Cancel at NBBO*, *non-HFT Order 1 tick from NBBO*, *non-HFT Cancel 1 tick from NBBO*, *non-HFT Order > 1 tick from NBBO*, and *non-HFT Cancel > 1 tick from NBBO*. The *HFT* named variables capture directional activity by

²⁴ Brogaard, Hendershott, and Riordan (2014) use a state-space model, similar to our VAR. The state-space model decomposes prices into permanent and transitory components and relates changes in both to trading variables. The state-space model is computationally and econometrically too complex to estimate on trade and quote variables as in our setup. The VAR allows for similar interpretation without explicitly modeling transitory price movements but including limit order submissions and cancellations.

HFT firms, and the *non-HFT* variables capture directional activity by non-HFT firms. Below we use “activity” and “directional activity” interchangeably. The *Trade* variables take the value of +1 for buyer-initiated trades, −1 for seller-initiated trades, and 0 otherwise. The *Order* variables take the value of +1 for bids placed at the NBB, −1 for offers placed at the NBO, and 0 otherwise. The *Order 1 tick from NBBO* variables take the value of +1 for bids placed at one cent from the NBB, −1 for offers placed at one cent from the NBO, and 0 otherwise. The *Order > 1 tick from NBBO* variables take the value of +1 for bids placed at greater than one cent from the NBB, −1 for offers placed at greater than one cent from the NBO, and 0 otherwise. For cancels, analogous definitions apply, with signs such that cancels at the bid take the value −1 and cancels at the offer take the value +1.

The observations include all displayed orders between 9:45 a.m. EST and 3:45 p.m. EST. To be included, a stock-day must have at least 20 nonzero observations of each variable. This eliminates 34 stock-days. The impulse response function (IRF) is orthogonalized and order-independent, and reports the forecasted midpoint return in basis points after a +1 (buy event for orders and trades, sell event for cancels). The IRF is cumulative over the following 20 events. The VAR is estimated for each stock-day. The averages of these stock-day IRF estimates are reported in basis points. For HFTs and non-HFTs a * (**) next to the coefficient indicates that on average the coefficient differs from zero and is statistically significant at the 5% (1%) level. In the “Difference” column, a * (**) next to the value indicates that the difference between the HFT and non-HFT coefficients is statistically significant at the 5% (1%) level. Throughout the paper, standard errors are clustered by stock-day to control for contemporaneous correlation across stocks and autocorrelation within stocks as in Thompson (2011).

To obtain IRFs we invert the VAR into its vector moving average representation. We calculate the IRFs for an unexpected buy message to measure its impact on returns. The return IRFs capture all types of messages’ price impact, often referred to as the contribution to price discovery or the information content of a message/order (event).

The VAR linearly models the dynamics between all order types and returns. The IRFs capture the impact that innovations in orders have on subsequent orders and returns. Modeling the dynamics enables proper measurement of the sequence of events. For example, suppose that some negative news arrives. HFTs immediately race to pick off stale bid orders. The first successful HFT picks off all mispriced buy limit orders with a single sell market order. Then other HFTs update their ask orders to quickly gain priority at new (lower) prices in the limit order book. The initial sell market order leads to an immediate price change by executing against all orders at the bid. The subsequent new orders at the ask also cause a price change. The return IRF aggregates immediate and predictable responses to the initial sell market order. Hence, both the immediate price change and the lagged price change are attributed to the market order. The price change due to the predictable limit orders is not

Table IV
Return Impulse Response Function

The table reports stock-day average return impulse response functions (IRFs) from a vector-autoregression (VAR) for the 15 non-cross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. There is one equation for each of the variables listed in the table, for each HFT and non-HFTs, and the midpoint NBBO midpoint return. The VAR is in event time, where every message is an observation. *Trade - Change Price* takes the value of +1 for buy-initiated trades (sell-initiated trades, -1, are defined analogously) that consume all depth at the NBO. *Trade - Same Price* takes the value of +1 for buy-initiated trades that do not consume all depth at the NBO. *Improving Order* takes the value of +1 (-1) for bids (offers) placed inside the NBB (NBO). *Order Placement at NBBO* takes the value of +1 for bids placed at the NBB, -1 for offers placed at the NBO, and 0 otherwise. *Order 1 tick from NBBO* takes the value of +1 for bids placed one cent from the NBB, -1 for offers placed one cent from the NBO, and 0 otherwise. *Order > 1 tick from NBBO* takes the value of +1 for bids placed more than one cent from the NBB, -1 for offers placed at more than one cent from the NBO, and 0 otherwise. For cancels, the analogous definitions apply with signs such that cancels at the bid take the value -1 and cancels at the offer take the value +1. Observations include all displayed messages between 9:45 a.m. EST and 3:45 p.m. EST. The IRF is cumulative over 20 events. For HFTs and non-HFTs, a * (**) next to the coefficient indicates that the coefficient is statistically different from zero at the 5% (1%) significance level using standard errors clustered by stock and by day. In the “Difference” column, * (**) next to the coefficient indicates that the HFTs and non-HFTs coefficients are statistically different from each other at the 5% (1%) significance level using standard errors clustered by stock and by day.

	HFT (1)	Non-HFT (2)	Difference (3)
Trade - Change Price	3.25**	2.99**	0.26**
Trade - Same Price	1.00**	0.72**	0.28**
Improving Order	1.56**	1.12**	0.44**
Worsening Cancel	2.00**	2.20**	-0.20**
Order Placement at NBBO	0.17**	0.14**	0.03**
Order Cancel at NBBO	0.01**	-0.03**	0.04**
Order 1 tick from NBBO	0.03**	0.04**	-0.01**
Cancel 1 tick from NBBO	0.01**	-0.01**	0.02**
Order > 1 tick from NBBO	-0.11**	-0.01**	-0.10**
Cancel > 1 tick from NBBO	0.02**	0.01**	0.01**

attributed to the limit order because the limit orders are not a “surprise” or innovation in the VAR system.

Table IV reports the estimates of the return IRFs from the VAR and can be interpreted as the average permanent price impact of an order type, aggressiveness, and participant combination. Orders generally have a positive permanent price impact. Trades that move the NBBO have the highest price impact, followed by limit orders and cancels that move the NBBO. Market orders that move the NBBO are limit orders priced aggressively with enough size to consume all liquidity at the best price.

Activity behind the best prices contributes relatively little to price discovery on average. HFT and non-HFT orders greater than one tick behind the NBBO both have a negative price impact. This is likely due to order submission geared

toward establishing limit order book priority. If the market is likely to tick in the direction of the submitted order, for example, the market moves up after a sell limit order is submitted, the order will have a negative price impact, but will also have time priority in the limit order book.²⁵ Under this interpretation, the order does not cause the price movement, rather it predicts the price movement. This seems to be fairly prevalent. It does not imply that randomly placed orders away from the inside have negative price impact, or even zero price impact. Rather, it implies that orders are placed away from the inside when prices are about to move toward the order, increasing its likelihood of execution, for example, price declines after a buy limit order is placed below the best bid price.

HFTs' directional activity generally moves prices more than non-HFTs' activity. Column (3) reports differences between HFT and non-HFT IRFs in basis points. Consistent with previous studies on HFTs and price discovery (Brogaard, Hendershott, and Riordan, 2014), we find that HFTs' market orders move prices more than the market orders of non-HFTs. Orders that move the NBBO also contribute considerably to price discovery. HFTs' limit order submissions that move the NBBO move prices by 1.56 basis points on average, whereas the same order for non-HFTs moves prices by 1.12 basis points. With the exception of cancellations that move the NBBO and activity behind the best prices, HFTs' orders impact prices more than do non-HFTs' orders.

It is unclear, however, what makes an average HFT order more informative than an average non-HFT order. Similar orders should have similar price impacts conditional on placement in the limit order book. Potential explanations are that HFTs use market data to learn when prices are likely to move (e.g., the depth at the best bid and ask prices is not equal), or may react faster to common information (e.g., the futures or ETF price moves). Both possibilities could lead to higher price impacts. HFTs learning from market data could lead them to trade more when price impacts are higher than average. HFTs reacting faster to common signals would lead their orders to have a higher price impact than slower subsequent orders from non-HFTs. We explore these possibilities below.

The VAR estimates used in Table IV also generate IRFs to explain how innovations in buy and sell messages predict the direction of subsequent buy and sell messages within and across order types and participant types. These IRFs provide evidence on the responses of HFTs and non-HFTs to market activity. As in Table IV, Table V reports the average stock-day IRF estimates.

The rows in Table V correspond to the message variable being shocked by one unit. The columns correspond to participant and message responses. For instance, the "Trade Δ NBBO HFT" row provides IRF estimates for how all HFT and non-HFT message types respond to an HFT buy trade that changes the NBBO. Table V presents the diagonal effects, as in Biais, Hillion, and Spatt (1995), which show that orders follow similar orders. Summing the columns

²⁵ Cancels by non-HFT at the NBBO and one tick behind the NBBO also have negative price impacts; Table VI explores possible explanations.

Table V
Message Impulse Response Function

The table reports stock-day average message impulse response functions (IRFs) from the vector-autoregressions (VAR) in Table IV for the 15 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. The VAR is in event time, where every message is an observation. The rows corresponds the variable being shocked by one unit. The columns represent the variable being affected. *Trade ΔNBO HFT (non-HFT)* takes the value of +1 for buy-initiated trades (sell-initiated trades, -1, are defined analogously) that consume all depth at the NBO. *Trade NBO HFT (non-HFT)* takes the value of +1 for buy-initiated trades that do not consume all depth at the NBO. *Order ΔNBO HFT (non-HFT)* takes the value of +1 (-1) for bids (offers) placed inside the NBB (NBO). *Order NBO HFT (non-HFT)* takes the value of +1 for bids placed at the NBB, -1 for offers placed at the NBO, and 0 otherwise. For cancels, the analogous definitions apply with signs such that cancels at the bid take the value -1 and cancels at the offer take the value +1. Observations include all displayed messages between 9:45 a.m. EST and 3:45 p.m. EST. The IRF is cumulative over 20 events. For HFTs and non-HFTs a * (***) next to the coefficient indicates that the coefficient differs from zero and is statistical significance at the 5 and 1% level, respectively using standard errors clustered by stock and by day. In the “Difference” column, * (***) next to the coefficient indicates that the HFT and non-HFT coefficients are statistically different from each other at the 5% (1%) significance level using standard errors clustered by stock and by day.

Variable	Trade		Order		Cancel		Trade		Order		Cancel		Trade		Order		Cancel		Trade		Order		Cancel	
	ΔNBO	HFT	ΔNBO	HFT	ΔNBO	HFT	ΔNBO	HFT	ΔNBO	HFT	ΔNBO	HFT	ΔNBO	HFT	ΔNBO	HFT	ΔNBO	HFT	ΔNBO	HFT	ΔNBO	HFT	ΔNBO	HFT
Trade ΔNBO HFT	0.00	0.08**																						
Trade NBO HFT	0.17**	1.20**																						
Order ΔNBO HFT	0.01**	0.05**																						
Cancel ΔNBO HFT	-0.01**	-0.02**																						
Order NBO HFT	0.01**	0.02**																						

(Continued)

Table V—Continued

[illegible]

across rows shows that HFTs' response to market activity is greater than that of non-HFTs. HFTs respond more to their own activity than to the activity of non-HFTs and more than non-HFTs respond to the activity of other non-HFTs. HFT activity predicts almost no additional activity from non-HFTs. That HFTs closely monitor market activity is consistent with the predictions of GPR in that speculators' order submission strategies are more sensitive to market activity than are non-speculators' order submission strategies.

Activity across market participants can be positively correlated for several reasons (Biais, Hillion, and Spatt, 1995). First, investors may react to common signals but with different speeds. Second, HFTs could be front-running non-HFTs' orders by detecting large non-HFT orders that are split over time or across exchanges (Hirschey, 2017; van Kervel and Menkveld, 2019; Korajczyk and Murphy, 2019). In both of these cases, slower investors will follow faster investors using the same order type. Third, investors may learn from one another. In an anonymous market, learning should be symmetric across investor types, for example, HFTs learn the same from HFTs' and non-HFTs' orders. The sequencing of order activity and its symmetry across HFTs and non-HFTs shed some light on how to disentangle learning from reacting to common signals with different speeds or quasi-front-running.

Table V shows that non-HFT activity mostly follows previous non-HFT activity but does not follow HFT activity. HFT activity follows both HFT and non-HFT activity but HFTs' orders do not generally anticipate non-HFTs' orders in the same direction. We find some evidence that HFTs may react more quickly to public information. For instance, HFT orders that move the NBBO negatively predict the same non-HFT orders and non-HFT market orders at the NBBO. While these types of problematic behavior do not represent the majority of HFTs' activity, this does not imply that HFTs never adversely select non-HFTs or quasi-front run their orders. HFT activity following the activity of HFTs and non-HFTs is also consistent with HFTs learning from other participants.

The largest difference in return IRFs between HFTs and non-HFTs from Table IV is for limit order submissions that move the NBBO (0.44 basis point difference). This could arise from HFTs' orders improving the NBBO by a greater amount or from HFTs' orders predicting future activity that will move the NBBO further. An intuitive way to compare the response to HFTs' orders that move the NBBO and the same orders for non-HFTs is to sum across the largest three response coefficients for each. For an HFT's order that moves the NBBO, this leads to the following (largest) responses: an HFT order that moves the NBBO (0.22 basis points), an HFT order at the NBBO (0.65 basis points), and a non-HFT order at the NBBO (0.18 basis points), for a total response of 1.05 basis points of additional activity. The same for a non-HFT leads to the following (largest) responses: an HFT order that moves the NBBO (0.37), a non-HFT trade at the NBBO (0.21), and a non-HFT order that moves the NBBO (0.39), for a total response of 0.97 basis points of additional activity. The difference is not large but may partially explain the larger price impacts of HFT orders that move the NBBO.

The linear system of the VAR limits the structure of interdependency in the system dynamics. To further understand the role of HFTs' and non-HFTs' orders in price discovery, we use a related approach that simply measures the price movement following an order. This price impact can be calculated for different message types conditional on some sources of public information, market conditions, and order sequences. Comparing the unconditional and conditional relative price impacts of HFTs' and non-HFTs' orders demonstrates the extent to which HFTs' higher price impact stems from things that can be thought of as public information.

Table VI reports averages of ordinary least squares (OLS) regressions of the 10-second signed price impacts (in basis points) performed for each stock-day on all messages that change or are at the NBBO. Signed price impacts capture the change in the midpoint in the direction of the order over the 10 seconds following the time the order hits the order book, and provide an intuitive way to understand how orders affect prices. For instance, if the midpoint for a \$10 stock increases (decreases) by \$0.10 after 10-seconds after the entry of a limit buy (sell) order, that order would be associated with a 100 basis point price impact. Midpoints that move in the opposite direction of placed orders have a negative price impact. Table IV shows that orders behind the best price in the order book have a low price impact. To simplify presentation and exposition, these orders are excluded from Table VI.

The first two columns of Table VI report the price impact results for "Trades." The next two columns report results for new "Orders" and the last two columns for "Order Cancels." *HFT* takes the value of 1 if an HFT placed the message, and zero otherwise. The variables in our analysis are as follows: *Order Size* is the number shares, *Stock Volatility* is the absolute value of the previous 10-second stock return (in percent), *XIU Volatility* is the absolute value of the previous 10-second return of the TSX 60 exchange-traded fund, XIU (in percent), *Relative Spread* is the stock's contemporaneous bid-ask spread divided by the midpoint price (in percent), *Lag Stock Return (1–10 second)* is the signed 10-second lagged stock return (in percent), *Lag XIU Return (1–10 second)* is the signed 10-second lagged XIU return (in percent), *Limit Order Book Imbalance* is the stock's contemporaneous signed limit order book imbalance defined as (depth at best bid price – depth at best ask price)/(depth at best bid price + depth at best ask price) for buy messages and the negative of this value for sell messages. In addition, we include *Hidden Order*, which plays two roles depending on whether the column is for trades or new orders²⁶: *Hidden Order* takes the value of 1 if the message was a hidden order in the "Order" column and or if the trade executes against a hidden order in the "Trade" column, and zero otherwise. A trade executing against a hidden order can be thought of as the trade being executed in different market conditions as hidden orders may be predictable. We next include *Change NBBO*, the immediate change (in basis points) in the NBBO due to the message. This variable controls for the possibility that HFTs

²⁶ Hidden orders represent a relatively small portion of orders submitted (5.27%), cancelled (3.15%), and executed (6.53%).

Table VI
Simple Price Impacts

The table reports averages of OLS regressions on the 10-second signed price impacts (in basis points) for each stock -day on all messages that change or are at the NBBO for the 15 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. Columns (1) and (2) are for Trades, columns (3) and (4) are for Orders, and columns (5) and (6) are for Order Cancels. *HFT* takes the value of 1 if an HFT placed the message, and 0 otherwise. Order Size is the number of shares in the order. *Stock Volatility* is the absolute value of the previous 10-second stock return (in percent). *XIU Volatility* is the absolute value of the previous 10-second return of the TSX 60 exchange-traded fund, XIU (in percent). *Relative Spread* is the stock's contemporaneous bid-ask spread divided by the midpoint price (in percent). *Lag Stock Return (1–10 second)* is the signed 10-second lagged stock return (in percent). *Lag XIU Return (1–10 second)* is the signed 10-second lagged XIU return (in percent). *Limit Order Book Imbalance* is the stock's contemporaneous signed limit order book imbalance defined as (depth at best bid price – depth at best ask price)/(depth at best bid price + depth at best ask price) for buy messages, and the negative of this value for sell messages. *Hidden Order* takes the value of 1 if the message was a hidden order (or for the Trade analysis, was executed against a hidden order), and 0 otherwise. *Change NBBO* is the immediate change (in basis points) in the NBBO because of the message. *Order Cancel followed by Order by Same Trader (Opposite Direction)* takes the value of 1 when the same trader immediately cancels a buy (offer) order followed by a buy (offer) order (or the reverse), and 0 otherwise. *Order Cancel Followed by Trade by Same Trader (Opposite Direction)* takes the value of 1 when the same trader immediately cancels a buy (sell) order followed by a buy (sell) trade (or the reverse), and 0 otherwise. * (**) next to the coefficient indicates that the coefficient is statistically different from zero at the 5% (1%) significance level using standard errors clustered by stock and by day.

	Trade (1)	Trade (2)	Order (3)	Order (4)	Order Cancel (5)	Order Cancel (6)
Intercept	2.38** (0.03)	1.79** (0.03)	0.69** (0.01)	0.62** (0.03)	–0.07** (0.01)	–0.02 (0.04)
HFT	1.58** (0.04)	0.96** (0.02)	0.16** (0.01)	0.11** (0.01)	0.15** (0.01)	0.10** (0.01)
Order Size		0.86** (0.06)		0.27** (0.01)		–0.06** (0.01)
Stock Volatility		–1.30** (0.20)		0.83** (0.09)		–0.04 (0.09)
XIU Volatility		–0.43* (0.18)		0.09* (0.04)		–0.02 (0.06)
Relative Spread		7.75** (0.43)		–1.74** (0.23)		0.39 (0.22)
Lag Stock Return (1–10 second)		–5.03** (0.37)		–6.28** (0.29)		–5.51** (0.30)
Lag XIU Return (1–10 second)		1.87** (0.20)		0.49** (0.08)		0.37** (0.06)
Limit Order Book Imbalance		0.50** (0.02)		0.24** (0.01)		0.29** (0.01)
Hidden Order		–1.38** (0.03)		–0.11** (0.01)		–0.12** (0.02)
Change NBBO		0.42** (0.01)		0.69** (0.01)		0.66** (0.01)
Order Cancel then Order by Same Trader (Opposite Direction)						–0.37** (0.01)
Order Cancel then Trade by Same Trader (Opposite Direction)						–0.99** (0.04)

or non-HFTs submit more aggressive orders or their trades are more likely to go through more levels of the order book. The remaining variables in this analysis include *Order Cancel followed by Order by Same Trader (Opposite Direction)*, which takes the value of 1 when the same trader immediately cancels a buy (sell) order followed by a buy (sell) order (or the reverse), and zero otherwise, and *Order Cancel Followed by Trade by Same Trader (Opposite Direction)*, which takes the value of 1 when the same trader immediately cancels a buy (sell) order followed by a buy (sell) trade (or the reverse), and zero otherwise.

That HFT limit orders that move the NBBO have a larger price impact than the same orders from non-HFTs could be due to from HFTs conditioning their order submission on public information or to submitting orders more when the price impact is high. For example, HFTs could be using the information in futures or ETF prices to trade in the constituent stocks. Alternatively, as shown in Brogaard, Hendershott, and Riordan (2014) for trades, HFT may submit more orders in the direction of limit order imbalances, which leads price impacts to be higher. The price impact regressions control for some obvious sources of public information and market conditions when price impacts may be high or low.

A notable result in Table VI is that controlling for public information and market conditions explains some of the HFTs' higher price impact.²⁷ For trades, the coefficient for HFTs decreases from 1.58 basis points when not conditioning on public information to 0.96 basis points when conditioning on public information and from 0.16 to 0.11 for new orders. This also holds for order cancels in that HFTs observe a statistically significantly lower price impact. The HFT coefficient falls by about one-third with the inclusion of the public information variables. One interpretation of this result is that one-third of HFTs' contribution to price discovery is due to public information. Significantly more analysis is needed, however, to gauge the accuracy of this intuitive estimate.

An additional result that follows from Biais, Hillion, and Spatt (1995) is that new orders submitted with wide spreads will have a lower price impact than the same orders submitted with tight spreads. Orders submitted within the spread when spreads are wide are more likely to be driven by a desire to be at the top of the queue than information motives. Conversely, trades submitted with wide spreads are more likely to be driven by informed traders that are willing to pay higher liquidity costs to profit from their information.

Price impacts are generally increasing in order size, lagged index returns, limit order book imbalances, and the immediate size of an NBBO change that price impacts decline with limit order book imbalances means that when a trader improves the NBBO with a new order on the side with more depth, the order appears less motivated by information. The trader likely believes that

²⁷ While controlling for public information explains some of HFTs' greater price impact, it is difficult to know how much of HFTs' price discovery would occur without HFTs. Estimating how non-HFTs would behave and incorporate information into prices without the presence of HFTs requires the exogenous removal or reduction of HFTs. Unfortunately, there is no such event in our sample period.

adverse selection is low and the new order is simply trying to get priority over the large depth at the previous NBBO. These orders have lower price impacts. When an order is placed at or better than the NBBO on the side with less depth, the order appears more motivated by information and has a larger price impact. Price impacts are generally decreasing in lagged signed stock returns and for hidden orders. Hidden orders seem to be submitted by uninformed market participants as they have lower price impacts (Zhu, 2013). Trades executing against hidden orders have smaller price impacts, likely because the midpoint of the bid-ask spread does not incorporate the hidden order.

Table IV reports negative IRFs for non-HFT cancels at and more than one tick behind the NBBO. The last two rows of Table VI report order sequences that include cancels that may lead to below average price impacts. For instance, a trader may decide to cancel an order because the NBBO has moved away from their order and then submit a new, more aggressively priced order. This would lead the initially cancelled order to have a below average price impact. The results show that cancels followed by orders and trades in the opposite direction do in fact lead to lower price impacts for cancels. Also, as expected, the coefficient is lower when the subsequent order is less aggressively priced (-0.37 vs. -0.99).

The previous results are informative about the average price impact of a message. However, given the autocorrelation and cross-correlation in message activity in Table V, one cannot simply weight the IRFs by the message frequencies in Table III to determine each message type's contribution to overall price discovery. The Hasbrouck (1991b) variance decomposition weights the IRFs by the variance of innovations in each order type to calculate the total contribution to price discovery by participant, order type, and relative aggressiveness in the order book. The VAR reported in Table IV produces the estimates to calculate the variance decomposition reported in Table VII.²⁸

In total, market orders at and that change the best price contribute roughly 30% to price discovery. Another 45% of price discovery occurs via limit orders. Of that 45%, 88% is impounded via limit orders that change the best price. In fact, 62% of price discovery is attributed to orders and trades that move the best price. This suggests that (1) the order book significantly contributes to price discovery, and (2) orders that change the best price contain more information than those that do not.

HFT limit orders that change the NBBO are more than twice as important as the same orders from non-HFTs (19.6% vs. 8.2%). This is explained largely by the fact that HFTs submit twice as many orders that change the NBBO than do non-HFTs (2.15% vs. 1.29% in Table III). The price impact of HFT orders that move the NBBO is larger than the same orders for non-HFTs, leading HFT

²⁸ While the variance decomposition contains every event that can cause prices to change, the linear structure of the VAR explains only roughly 75% of the variance of the efficient price; the remainder is variance of the error term in the return equation. To ensure that our variance decomposition across order types sums to 100%, we use the amount of the efficient price that is explained as the denominator when calculating each order type's contribution to price discovery.

Table VII
Variance Decomposition

The table reports the stock-day average variance decomposition from the vector-autoregression (VAR) used in Table IV for the 15 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013.

	HFT (1)	Non-HFT (2)
Trade - Change Price	8.8%	11.6%
Trade - Same Price	3.0%	6.8%
Improving Order	19.6%	8.2%
Worsening Cancel	8.8%	5.2%
At NBBO Order	1.5%	1.0%
At NBBO Cancel	0.2%	0.1%
Order 1 tick from NBBO	0.1%	0.04%
Cancel 1 tick from NBBO	0.1%	0.04%
Order > 1 tick from NBBO	0.1%	0.04%
Cancel > 1 tick from NBBO	0.1%	0.05%

limit orders that move the NBBO to contribute more to price discovery than any other participant/order-aggressiveness combination.

Trades that move the NBBO and cancels that move the NBBO are also important in terms of price discovery. The differences and magnitudes in price discovery are commensurate with their relative frequencies and price impacts. For instance, non-HFT trades that move the NBBO are more frequent than those of HFTs and have a higher price impact. This leads their overall role in price discovery to be greater than for HFT trades that move the NBBO. For cancels, HFTs are roughly 2.5 times more likely to cancel orders that move the NBBO than are non-HFTs. This leads their role in price discovery to be higher (8.8%) than for non-HFTs (5.2%) despite the slightly lower price impact of HFT cancels.

HFTs contribute the most to price discovery via their limit orders and non-HFTs via their market orders. This suggests that HFTs make different order submission decisions conditional on their information than non-HFTs. This result is consistent with the finding in theoretical literature (GPR and others) that speculators use limit orders more.

A testable prediction from the literature on limit order books is that speculators reduce their submission of limit orders as volatility increases (Bloomfield, O'Hara, and Saar (2005) and GPR). Table VIII directly examines the relationship between stock volatility and message frequency. Table VIII reports results from OLS regressions on message frequencies by limit/market order and HFT/non-HFT from Table III. In this analysis, % by Limit Order is the percentage of messages by limit order book submissions and cancellations, % by HFT Limit Order is the percentage of limit order book submissions and cancellation by HFTs, % by HFT Limit Order / % HFT is the percentage of HFT limit orders and cancellations by total HFT, % by HFT Limit Order / % by Limit Order is the percentage of HFT limit orders normalized by the percentage of

Table VIII

Stock Volatility and Message Frequency by HFT and Non-HFT and by Order Type

The table reports coefficients from pooled OLS regressions on daily message frequencies (as Table III) for the 15 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. *% by Limit Order* is the percentage of messages by limit order book submissions and cancellations. *% by HFT Limit Order* is the percentage of limit order book submissions and cancellations by HFTs. *% HFT by Limit Order / % HFT* is the percentage of HFT limit orders and cancellations by total HFTs. *% HFT by Limit Order / % by Limit Order* is the *% HFT* limit orders normalized by the percentage of total limit orders. *% by HFT Limit Order / % by Market Order* is the percentage of HFT limit orders normalized by the percentage of market orders. *% by non-HFT Limit Order* is the percentage of non-HFT limit orders. *% by non-HFT Limit Order / % non-HFT* is the daily absolute return from the previous trading day, in percent. *Market Cap* is the market capitalization of the stock, in billions, and *1 / Price* is the reciprocal of the stock price. * (**) next to the coefficient indicates that the coefficient is statistically different from zero at the 5% (1%) significance level using standard errors clustered by stock and by day.

	<i>% by Limit Order</i> (1)	<i>% by HFT Limit Order</i> (2)	<i>% by HFT Limit Order / % HFT</i> (3)	<i>% by HFT Limit Order / % by Limit Order</i> (4)	<i>% by HFT Market Order / % Market Order</i> (5)	<i>% by Non-HFT Limit Order</i> (6)	<i>% by Non-HFT Order / % Non-HFT</i> (7)
Intercept	95.59** (0.17)	43.68** (0.81)	96.43** (0.19)	45.37** (0.82)	25.43** (0.64)	51.92** (0.75)	93.72** (0.31)
Stock Volatility (<i>t</i> − 1)	−0.99** (0.16)	−3.40** (0.67)	−0.85** (0.16)	−3.08** (0.73)	1.08* (0.48)	2.41** (0.68)	−0.40 (0.34)
Market Cap	−0.11** (0.01)	0.62** (0.03)	0.05** (0.01)	0.73** (0.03)	−0.23** (0.02)	−0.72** (0.03)	−0.36** (0.02)
1/Price	−52.31** (1.38)	25.15** (4.63)	−8.90** (1.92)	58.99** (4.75)	−36.67** (5.36)	−77.46** (3.93)	−96.72** (2.24)

total limit orders, % by *HFT Limit Order* / % by *Market Order* is the percentage of HFT limit orders normalized by the percentage of market orders, % by *non-HFT Limit Order* is the percentage of non-HFT limit orders, and % by *non-HFT Limit Order* / % *non-HFT* is the percentage of non-HFT limit orders normalized by the percentage of total non-HFT orders. Observations are at the stock-day level. Because stock volatility is persistent and to avoid simultaneity between volatility and traders' order choice, *Stock Volatility* ($t - 1$) is the daily absolute return in percentage the prior day. To focus on within-stock volatility, we include two cross-sectional controls²⁹: *Market Capitalization* is the market capitalization of the stock in billions and $1/Price$ is the reciprocal of the stock price.

The predictions of Bloomfield, O'Hara, and Saar (2005) and GPR for informed traders are generally confirmed in Table VIII. For example, the frequency of limit order submissions falls by 0.99% per 1% increase in the previous day's volatility. The decrease in limit order submissions is driven by the reduction in HFT limit orders. A 1% increase in previous-day stock volatility leads HFT limit order frequency to fall by a statistically significant 3.08% and 3.40%, relative to limit orders and overall activity, respectively. A 1% increase in volatility leads non-HFT limit order submissions to increase 2.41% relative to overall activity. Overall limit order submissions are less frequent for larger stocks and for lower priced stocks. However, HFT limit order submissions are more frequent in larger stocks and lower priced stocks. Table VIII suggests that HFT and non-HFT limit order submission frequencies vary differently with volatility. Table IX tests whether the importance of these types of orders for price discovery changes with volatility.

Table IX examines the relationship between stock volatility and contribution to price discovery. The table reports results from OLS regressions on different components of the variance decomposition from Table VII. The various price discovery ratios follow the same logic as the ratios used for message frequencies in Table VIII. Observations are at the stock-day level and volatility, market capitalization, and inverse price remain the same.

Mechanically, the numerator in the variance decomposition is the product of the variance of the surprise/innovation of the order flow and the squared price impact of an order (IRF). This suggests that if a certain type of order is submitted less frequently (or less unexpectedly) then the contribution to price discovery by that type of order should fall. However, average price impacts and the innovation component of orders may both vary with volatility. Consistent with theory, Table IX shows that limit orders contribute less to price discovery as volatility increases. The contribution of limit orders to price discovery falls by 1.04% per 1% increase in previous-day volatility. Also consistent with the frequency results of Table VIII, we find that the decrease in the explanatory power of limit orders is driven by a decrease in HFT limit order contribution to price discovery. For example, HFT limit orders make up 1.31% less of the total

²⁹ Internet Appendix Tables AVII and AVIII replace the price and firm size controls with firm fixed effects. The volatility results are similar.

Price Discovery without Trading

1647

Table IX
Stock Volatility and Percentage of Price Discovery by HFT and Non-HFT and by Order Type

The table reports the coefficient from pooled OLS regressions on the different components of the daily variance decomposition (from Table VII) for the 115 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. % by *Limit Order* is the percentage of the efficient variance explained by limit orders (cancellations and submissions) normalized by the fraction of explained variance from the variance decomposition. % by *HFT Limit Order* is the percentage of the variance decomposition estimates explained by HFT limit orders normalized by the explained variance % by *HFT Limit Order* / % *HFT* is the percentage of the variance decomposition estimates explained by HFT limit orders normalized by the percent HFT explained. % by *HFT Limit Order* / % by *Limit Order* is the percentage of the variance decomposition estimates explained by HFT limit orders normalized by the percentage explained by limit orders. % by *HFT Limit Order* / % by *Market Order* is the percentage of the variance decomposition estimates explained by HFT limit orders normalized by the percentage explained by Market orders. % by *non-HFT Limit Order* is the percentage of the variance decomposition estimates explained by non-HFT limit orders normalized by the explained variance. % by *non-HFT Limit Order* / % *non-HFT* is the percentage of the variance decomposition estimates explained by non-HFT limit orders normalized by the percentage explained by non-HFT. Observations are at the stock-day level. *Stock Volatility* ($t - 1$) is the daily absolute return from the previous trading day, in percent. *Market Cap* is the market capitalization of the stock, in billions, and $1/Price$ is the reciprocal of the stock price. * (***) next to the coefficient indicates that the coefficient is statistically different from zero at the 5% (1% significance level using standard errors clustered by stock and by day).

	% by Limit Order (1)	% by HFT Limit Order (2)	% by HFT Limit Order/ % HFT (3)	% by HFT Limit Order/ % by Limit Order (4)	% by HFT Market Order/ % Market Order (5)	% by Non-HFT Limit Order (6)	% by Non-HFT Limit Order/ % Non-HFT (7)
Intercept	60.48** (0.91)	20.68** (1.03)	54.19** (1.29)	33.65** (1.17)	48.21** (0.86)	39.79** (0.67)	71.17** (1.00)
Stock Volatility ($t-1$)	-1.04** (0.26)	-1.31** (0.36)	-0.80* (0.36)	-1.10* (0.46)	-0.66 (0.32)	0.27 (0.27)	-0.51 (0.39)
Market Cap	0.89** (0.28)	8.86** (0.28)	7.13** (0.37)	13.83** (0.38)	-3.04** (0.31)	-7.96** (0.26)	-10.68** (0.38)
1/Price	-33.86** (5.04)	19.25** (5.92)	19.34* (7.16)	65.55** (6.52)	-36.02** (7.24)	-53.10** (3.35)	-86.66** (6.12)

variance explained per 1% increase in previous-day volatility. Thus, the contribution to price discovery by non-HFTs does not appear to vary with volatility.

III. Activity and Price Discovery within and across Markets

The market-wide price discovery analysis in Tables IV to VII shows that HFTs' limit orders are the primary channel through which price discovery occurs. In this section we begin by studying HFT price discovery across exchanges. We first ask whether HFTs' role in price discovery is concentrated on the new exchanges or is similar across the three largest exchanges. Next, we study potential cross-exchange sources of information by extending Tables IV to VII to examine how HFTs and non-HFTs respond to trading and orders within and across markets.³⁰ Finally, we examine directly how price changes migrate across exchanges.

Table X decomposes order activity by participant type and order type relative to the NBBO for each exchange. Most activity is concentrated at the best bid and offer (BBO) across all three exchanges.

Exchange 2 is the largest of the three Canadian exchanges. The activity breakdown in terms of order type participant frequencies across exchanges is comparable with the exception of exchange 3, where HFTs are responsible for 75% of activity versus 48% and 41% for exchanges 1 and 2, respectively. That HFTs are more active on smaller exchanges is consistent with Menkveld (2013) highlighting the importance of HFTs for smaller and newer exchanges. Roughly half of the activity is better than or at the best prices for all three exchanges.

Tables IV and V show that HFTs' new limit orders are the primary channel for price discovery. The analyses are in event time, which does not account for possibly very small time differences between HFTs and non-HFTs incorporating new information. For example, if a public information event, say the S&P 500 futures price (or the TSX 60 ETF price) increasing, leads to both HFTs' and non-HFTs' buy limit order submissions with HFTs being slightly faster, event time analysis will attribute all price discovery to HFTs. A standard approach to measuring price discovery in calendar time is to use Hasbrouck's (1995) information shares, by examining quotes in different markets for the same security (e.g., Huang, 2002). This approach can be directly extended to the best quotes by different market participants. The information shares approach decomposes variation in the common efficient price into individual components attributable to specific markets or participants.

The information share methodology focuses on innovations in different groups' prices (quote midpoints). The information share of a group is measured as that group's contribution to the total variance of the common (random-walk) component. We calculate the price path for each group (HFT/non-HFT or Exchange), where the price vector p_t represents the prevailing prices for each group i and is given as $p_t^i = m_t + \epsilon_t^i$. Prices are assumed covariance-stationary.

³⁰ As in the Table II, Internet Appendix Table AIX provides descriptive statistics for each exchange.

Table X
Message Frequency by Exchange

The table reports the average stock-day distribution of messages by HFTs and non-HFTs on all three exchanges for the 15 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. The table is similar to Table III except that each message type is further broken down by exchange 1, 2, or 3.

	HFT (1)	Non-HFT (2)
Exchange 1		
Trade - Change Price	0.23%	0.29%
Trade - Same Price	0.56%	2.48%
Improving Order	1.16%	1.13%
Worsening Cancel	0.62%	0.97%
Order Placement at NBBO	14.93%	8.46%
Order Cancel at NBBO	11.65%	5.29%
Order 1 tick from NBBO	3.99%	1.39%
Cancel 1 tick from NBBO	4.40%	2.30%
Order > 1 tick from NBBO	4.77%	12.80%
Cancel > 1 tick from NBBO	6.00%	16.57%
Total Number of Observations	16,250	
Exchange 2		
Trade - Change Price	0.28%	0.35%
Trade - Same Price	1.11%	4.67%
Improving Order	2.12%	1.70%
Worsening Cancel	0.73%	0.72%
Order Placement at NBBO	12.79%	11.13%
Order Cancel at NBBO	10.39%	6.98%
Order 1 tick from NBBO	2.12%	2.02%
Cancel 1 tick from NBBO	3.16%	3.36%
Order > 1 tick from NBBO	3.89%	11.77%
Cancel > 1 tick from NBBO	4.87%	15.85%
Total Number of Observations	33,669	
Exchange 3		
Trade - Change Price	0.09%	0.17%
Trade - Same Price	0.37%	2.18%
Improving Order	2.90%	0.79%
Worsening Cancel	1.34%	0.25%
Order Placement at NBBO	17.95%	4.43%
Order Cancel at NBBO	16.00%	3.27%
Order 1 tick from NBBO	6.32%	0.95%
Cancel 1 tick from NBBO	6.92%	1.56%
Order > 1 tick from NBBO	11.12%	4.96%
Cancel > 1 tick from NBBO	11.96%	6.46%
Total Number of Observations	22,547	

The common efficient price path is a random walk process, $m_t = m_{t-1} + u_t$, where $E(u_t) = 0$, $E(u_t^2) = \sigma_u^2$, and $E(u_t u_s) = 0$ for $t \neq s$. The price process vector can be modeled as a VMA:

$$\Delta p_t = \varepsilon_t + \psi_1 \varepsilon_{t-1} + \psi_2 \varepsilon_{t-2} \dots,$$

where ϵ is a vector of innovations with zero mean and a variance matrix of Ω , and Ψ represents the polynomial in the lag operator. The information share is:

$$InfoShare_j = \frac{\Psi_j^2 \Omega_j}{\Psi \Omega \Psi'},$$

where $InfoShare_j$ is the fraction of price discovery attributable to participant (or exchange) j , the numerator is the variance of the efficient price attributable to participant j , and the denominator is the total variance of the efficient price. As discussed in Hasbrouck (1995), when multiple series move at the same time, the information share cannot be uniquely attributed to any series. In our setting this occurs if both HFTs' and non-HFTs' prices move at the same time. Information shares are typically estimated at a fixed sampling frequency. Higher sampling frequencies allow price discovery to be more uniquely attributed to HFTs and non-HFTs, but also attribute price discovery occurring close together to the faster participant group, presumably HFTs. To balance this trade-off, we estimate information shares at the one-second frequency.

Table XI reports the information shares for HFTs and non-HFTs. As with the VARs in Tables IV and V, information shares are estimated for each stock-day. Table XI provides the average maximum and minimum information shares. Table XI further reports whether the average minimum HFT information share is statistically significantly greater than the average maximum non-HFT information share. Testing the minimum HFT information share against the maximum non-HFT information share indicates whether the HFT limit order price discovery results in Tables IV, VI, and VII, are simply due to HFTs updating their quotes faster than non-HFTs.

Table XI, Panel A, shows that the average 59% minimum information share for HFTs is statistically significantly greater than the 41% maximum value for non-HFTs. This result demonstrates that even with conservative timing assumptions, HFT limit order activity contributes more to price discovery than non-HFT activity. The HFT maximum information share attributes all of the common innovations in price discovery within the same second to HFTs and shows that HFTs account for 81% of price discovery. Thus, 22% of the price discovery occurs within the same second for HFTs and non-HFTs.

Panels B and C of Table XI extend the aggregate information share price discovery analysis in Panel A to prices from the different exchanges. Panel B calculates information shares by exchange. All exchanges contribute to price discovery. Exchange 1 has the highest minimum and maximum information shares. The wide gap between the maximum and minimum exchange information shares suggests that price discovery is well integrated across markets, as significant common price discovery occurs across exchanges within one second.

Panel C of Table XI breaks the information shares down by exchange and participant type. HFTs are responsible for at least 56% of the information on every exchange and at most 93% on exchange 3. HFTs are responsible for the least price discovery on the largest exchange, although still responsible for more price discovery on that exchange than non-HFTs. The results of

Table XI
Information Shares

The table reports the average stock-day Hasbrouck (1995) minimum and maximum information shares for the 15 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. Panel A evaluates HFT and non-HFT quotes. Panel B evaluates quotes on each exchange. Panel C evaluates HFT and non-HFT quotes on each exchange. * (**) next to the HFT-Min point estimate indicates that the difference between the HFT Min and the non-HFT Max estimates statistically significant at the 5% (1%) level using standard errors clustered by stock and by day.

	Min (1)	Max (2)
Panel A: By HFT		
HFT	0.59**	0.81
Non-HFT	0.19	0.41
Table B: By Exchange		
Exchange 1	0.29	0.58
Exchange 2	0.23	0.47
Exchange 3	0.14	0.36
Panel C: By HFT and Exchange		
Exchange 1		
HFT	0.75**	0.87
Non-HFT	0.13	0.25
Exchange 2		
HFT	0.56**	0.76
Non-HFT	0.24	0.44
Exchange 3		
HFT	0.84**	0.93
Non-HFT	0.07	0.16

Table XI suggest that HFTs’ information advantages are not driven primarily by differences across exchanges, for example, exchange-specific HFT friendly order types. Table XI also suggests that HFTs’ orders moving prices more than the orders of non-HFTs is explained not only by their being marginally faster than non-HFTs, as these information shares can mitigate very small speed advantages.

We next study potential cross-exchange sources of information. A source of HFT information could be activity in other markets. In Table V we show that HFTs are more responsive to overall activity than are non-HFTs. Table XII extends the VAR in Tables IV and V to capture order activity by HFTs and non-HFTs on the same and different exchanges. This is done by indexing all trade and order variables by exchange. Thus, there are 10 order variables on each of the three exchanges for both HFTs and non-HFTs, resulting in a system of 61 equations and variables. As before, we estimate the VAR for each stock-day and report stock-day averages.

Table XII
Message Impulse Response Function by Exchange

The table reports stock-day average message impulse response functions (IRF) from a vector-autoregression (VAR) similar to that in Table IV, except that each message is broken down by exchange 1, 2, or 3 for the 15 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. The rows correspond to the variable being shocked by one unit. The columns correspond to the variable being affected. Panel A reports the average IRFs for the same exchange (e.g., the IRF of an Order HFT innovation on Exchange 3 on a Cancel HFT on Exchange 3). Panel B reports the average IRFs for the other exchanges (e.g., the IRF of an Order HFT innovation on Exchange 3 on a Cancel HFT on Exchanges 1 and 2). * (**) next to the coefficient indicates that the coefficient is statistically different from zero at the 5% (1%) significance level using standard errors clustered by stock and by day.

Panel A: Same Exchange																								
Variable	Trade ΔNBBO HFT	(1)	Trade NBBO HFT	(2)	Order ΔNBBO HFT	(3)	Cancel ΔNBBO HFT	(4)	Order NBBO HFT	(5)	Cancel NBBO HFT	(6)	Trade ΔNBBO Non-HFT	(7)	Trade NBBO Non-HFT	(8)	Order ΔNBBO Non-HFT	(9)	Cancel ΔNBBO Non-HFT	(10)	Order NBBO Non-HFT	(11)	Cancel NBBO Non-HFT	(12)
Trade ΔNBBO HFT	0.00*		0.01**		0.28**		-0.00*		0.44**		-0.06**		-0.01**		-0.03**		0.01**		0.00		0.04**		0.02**	
Trade NBBO HFT	0.09**		0.54**		0.19**		-0.00**		0.30**		-0.11**		-0.00**		-0.03**		0.09**		0.00		0.03**		-0.00**	
Order ΔNBBO HFT	-0.00**		0.00**		0.22**		-0.10**		0.19**		-0.01**		-0.01**		-0.03**		0.00*		0.00**		0.02**		0.03**	
Cancel ΔNBBO HFT	-0.00**		-0.01**		-0.03**		0.02**		0.15**		0.04**		-0.00**		-0.03**		-0.02**		0.00		-0.01**		0.02**	
Order NBBO HFT	0.00**		0.01**		0.02**		-0.01**		0.23**		-0.27**		-0.00		-0.01**		-0.00		0.01**		0.03**		0.01**	
Cancel NBBO HFT	0.00**		0.00**		-0.05**		0.02**		-0.11**		0.14**		-0.00**		-0.01**		-0.00**		0.00**		0.01**		0.01**	
Trade ΔNBBO Non-HFT	-0.00		-0.01**		0.16**		0.00		0.40**		0.04**		-0.00**		0.05**		0.02**		0.00**		0.02**		0.00	
Trade NBBO Non-HFT	0.01**		0.05**		0.10**		-0.00**		0.30**		0.01**		0.03**		0.49**		0.01**		0.00**		0.03**		-0.00**	
Order ΔNBBO Non-HFT	-0.00**		0.00**		0.23**		-0.10**		0.19**		-0.01**		-0.01**		-0.03**		0.00*		0.00**		0.02**		0.03**	
Cancel ΔNBBO Non-HFT	-0.00**		-0.01**		-0.03**		0.02**		-0.15**		0.04**		-0.00**		-0.03**		-0.02**		0.00		-0.01**		0.02**	
Order NBBO Non-HFT	0.00**		0.00**		0.01**		0.01**		0.09**		0.00		0.00**		-0.00		0.01**		-0.01**		0.12**		-0.30**	
Cancel NBBO Non-HFT	0.00**		-0.00**		0.00*		0.00**		0.01**		0.05**		-0.00**		-0.03**		-0.03**		0.01**		-0.19**		0.10**	

(Continued)

Table XII—Continued

Variable	Panel B: Other Exchange											
	Trade ΔNBBO HFT (1)	Trade NBBO HFT (2)	Order ΔNBBO HFT (3)	Cancel ΔNBBO HFT (4)	Order NBBO HFT (5)	Cancel NBBO HFT (6)	Trade ΔNBBO Non-HFT (7)	Trade NBBO Non-HFT (8)	Order ΔNBBO Non-HFT (9)	Cancel ΔNBBO Non-HFT (10)	Order NBBO Non-HFT (11)	Cancel NBBO Non-HFT (12)
Trade ΔNBBO HFT	0.00**	-0.01**	0.05**	0.00*	0.17**	0.02**	-0.00	-0.01**	-0.01**	0.01**	0.02**	0.02**
Trade NBBO HFT	0.05**	0.14**	0.07**	0.01**	0.17**	-0.01*	0.01**	0.06**	0.00**	0.01**	0.03**	0.00**
Order ΔNBBO HFT	0.01**	0.02**	0.02**	0.00**	0.14**	0.03**	0.01**	0.03**	0.01**	-0.00	0.04**	0.02**
Cancel ΔNBBO HFT	-0.00**	-0.01**	-0.05**	0.01**	-0.05**	0.05**	-0.00**	-0.02**	-0.02**	0.01**	-0.04**	0.03**
Order NBBO HFT	0.01**	0.01**	0.03**	0.00**	0.15**	-0.01**	0.01**	0.02**	0.01**	0.00**	0.07**	-0.02**
Cancel NBBO HFT	-0.00**	-0.00**	-0.02**	0.02**	-0.00	0.13**	-0.00**	0.00**	-0.01**	0.01**	-0.01**	0.03**
Trade ΔNBBO Non-HFT	0.00**	-0.00**	0.05**	0.00	0.17**	0.04**	0.00**	0.00	0.01**	0.00**	0.01**	0.03**
Trade NBBO Non-HFT	0.01**	0.03**	0.04**	0.01**	0.16**	0.06**	0.02**	0.16**	0.01**	0.00**	0.04**	0.02**
Order ΔNBBO Non-HFT	0.01**	0.02**	0.02**	0.00**	0.14**	0.03**	0.01**	0.03**	0.01**	-0.00	0.04**	0.02**
Cancel ΔNBBO Non-HFT	-0.00**	-0.01**	-0.05**	0.01**	-0.05**	0.05**	-0.00**	-0.02**	-0.02**	0.01**	-0.04**	0.03**
Order NBBO Non-HFT	0.00**	0.01**	0.01**	0.01**	0.06**	0.01**	0.00**	0.01**	0.01**	0.01**	0.14**	-0.03**
Cancel NBBO Non-HFT	-0.00**	-0.00**	-0.00**	0.00**	0.02**	0.03**	-0.00**	-0.01**	-0.01**	0.01**	-0.06**	0.08**

Table XII provides the trade and order IRFs from the above-described VAR.³¹ To avoid presenting all of the possible cross-exchange IRFs, we do not report the coefficients on variables that are one tick or more away from the NBBO. Also, we group together IRFs for same and other exchanges. The same-exchange IRFs are the cross-exchange average response of each variable to the variables on that exchange. For example, the same-exchange HFT trade response to an HFT trade is the average of the IRFs for an HFT trade on exchange 1 response to an HFT trade on exchange 1, an HFT trade on exchange 2 response to an HFT trade on exchange 2, and an HFT trade on exchange 3 response to an HFT trade on exchange 3. The other-exchange IRFs are the cross-exchange average response of each variable to the variables on other exchanges. For example, the other-exchange HFT trade response to an HFT trade is the average of the IRFs for an HFT trade on exchange 1 response to an HFT trade on exchanges 2 and 3, an HFT trade on exchange 2 response to an HFT trade on exchanges 1 and 3, and an HFT trade on exchange 3 response to an HFT trade on exchanges 1 and 2. Panel A reports the average IRFs for the same exchange, while Panel B reports the average IRFs for the other exchanges.

The results presented in Panel A of Table XII are consistent with those presented in Table V. Similar to Table V, summing the coefficients for HFT events and response combinations provides insights into how HFT and non-HFT activity precedes and responds to particular events. For instance, summing the upper-left quadrant of Table XII, Panel A shows that HFTs' response to HFT activity is 3.65 basis points. Summing the absolute value of the lower-right quadrant shows that non-HFTs response to non-HFT activity is 1.62 basis points.

There is a strong diagonal effect, with HFTs responding more to market activity, in the same direction, than non-HFTs. There are differences in how the same- and other-market activity precedes HFT and non-HFT activity in the same direction. For HFTs, same-market activity precedes more same-market activity in the same direction than does other-market activity. For non-HFTs, other-market activity precedes more activity in the same direction than does same-market activity. This could arise from some non-HFTs' algorithms splitting large orders across markets.

Finally, we focus directly on how price changes migrate across exchanges. Prices changing across multiple exchanges are notable events when examining the role of HFTs' and non-HFTs' orders in cross-market price discovery. While these are only a small portion of activity, they allow for straightforward study of which orders initiate and finish the cross-market price changes. Table XIII examines message sequences where exchanges 2 and 3 start with the same best bid or offer price and both exchanges revise their bid/offer in the same direction. We focus on exchanges 2 and 3 in this analysis as these are the two largest

³¹ As in Table IX, Internet Appendix Table AX provides IRFs for each exchange.

Table XIII
Cross Exchange Price Change Sequences

The table reports the frequencies of price change sequences on Exchange 2 and Exchange 3 for the 15 noncross-listed stocks in the TSX 60. The sample period is from October 15, 2012 to June 28, 2013. The sequences capture all events in which Exchanges 2 and 3 start with the same bid or ask price, followed by one of the exchange's best bid or ask price changing, and subsequently the other exchange updating its quote in the same direction. The deviation and resolution must persist at least one millisecond. The statistics are calculated using data from October 15, 2012 to June 28, 2013. This analysis evaluates the joint probability of an HFT/non-HFT and Trade/Order/Cancel opening a sequence and an HFT/non-HFT and Trade/Order/Cancel subsequently updating the second exchange's quotes. The average stock-day has 1,021 sequences that persist for at least one millisecond.

Conditional on Open/Close Order Type and HFT, ≥ 1 Millisecond								
First	Second						% Open by Trader-Order (7)	% Open by Trader (8)
	HFT			Non-HFT				
	Trade (1)	Order (2)	Cancel (3)	Trade (4)	Order (5)	Cancel (6)		
HFT								
Trade	0.7%	—	0.9%	0.4%	—	0.2%	2.2%	
Order	—	27.1%	—	—	9.1%	—	36.1%	68.9%
Cancel	2.0%	—	17.7%	3.8%	—	7.0%	30.6%	
Non-HFT								
Trade	0.9%	—	3.0%	2.6%	—	2.0%	8.6%	
Order	—	6.9%	—	—	5.0%	—	11.9%	31.1%
Cancel	0.5%	—	4.6%	1.5%	—	4.1%	10.6%	
% Close by Trader- Order	4.2%	34.0%	26.2%	8.3%	14.1%	13.3%		
% Close by Trader		64.3%			35.7%			

exchanges by volume.³² The initial message that causes the two exchanges to go from the same price to different prices is referred to as the first message and the subsequent message on the second exchange that causes its price to move is referred to as the second message. Table XIII provides insights into the sequence of messages that result in common price discovery across the two exchanges. The first message does not need to immediately precede the second message. Table XIII reports the frequencies of message sequences, where the price changes on the two exchanges do not occur simultaneously (not in the same millisecond).

³² Internet Appendix Tables AXI and AXII repeat the analysis for exchanges 1 and 3 and exchanges 1 and 2. Internet Appendix Tables AXIII provides results on how long it takes price-change events to occur across the two markets.

Because Table XIII examines cases in which both exchanges' prices move in the same direction, not all message sequences are possible. For example, after a trade consumes all of the liquidity available at the best price on one exchange, the price on the same side of the market can change in the same direction on the other exchange only due to a trade or an order cancellation. Because there are fewer sequence possibilities, Table XIII reports joint frequencies/probabilities.

The results show that HFTs are responsible for initiating price changes that cross exchanges more often than non-HFTs. HFTs are also responsible for closing these price changes. The result that HFTs facilitate price discovery across markets is consistent with previous results on HFTs' important role in price discovery.

IV. Conclusion

The traditional view of price discovery is that trades reveal investors' private information while market makers' quotes reflect public information. Most stock exchanges and financial markets are now limit order books without designated market makers. The market/limit order choice by informed and uninformed investors determines the nature of price discovery and adverse selection.

Using regulatory data identifying limit orders and trades by HFTs and non-HFTs, we study the role of these two types of trades in price discovery and test theoretical predictions of models of limit order books. HFTs' market orders play a small role in price discovery while HFTs' limit orders play a large role. The widely stated concern that HFTs' market orders adversely select non-HFTs' limit orders plays a small role in price discovery. In addition, we find little evidence that HFTs use their speed advantage to front-run non-HFTs' orders. While these types of problematic behavior do not represent the majority of HFTs' activity, this does not imply that HFTs never adversely select non-HFTs or quasi-front-run their orders.

Consistent with theoretical models of limit order books, we find that price discovery switches from limit to market orders when volatility increases and that this change is due to HFTs. Volatility typically decreases market stability and HFTs' change in behavior could exacerbate the impact of volatility. A natural regulatory response to increase market stability could be to restrict HFTs' ability to demand liquidity when volatility increases. However, this could interfere with HFTs' risk management practices. In addition, identifying and regulating specific market participants is challenging. Incentives, such as reduced trading fees or privileges, are often needed for traders to self-identify. How to induce traders to produce the socially optimal level of market stability by continuing to provide liquidity via limit orders in stressful market conditions is an important topic for future research.

REFERENCES

- Ahn, Hee-Joon, Kee-Hong Bae, and Kalok Chan, 2001, Limit orders, depth and volatility: Evidence from the stock exchange of Hong Kong, *Journal of Finance* 56, 767–788.
- Anand, Amber, Sugato Chakravarty, and Terrence Martell, 2005, Empirical evidence on the evolution of liquidity: Choice of market versus limit orders by informed and uninformed traders, *Journal of Financial Markets* 8, 288–308.
- Bacidore, Jeffrey, and George Sofianos, 2002, Liquidity provision and specialist trading in NYSE-listed non-U.S. stocks, *Journal of Financial Economics* 63, 133–158.
- Baldauf, Markus, and Joshua Mollner, 2016, Fast traders make a quick buck: The role of speed in liquidity provision, Working paper, American Economic Association Annual Meeting.
- Biais, Bruno, Pierre Hillion, and Chester Spatt, 1995, An empirical analysis of the limit order book and the order flow in the Paris Bourse, *Journal of Finance* 50, 1655–1689.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas, 2015, Equilibrium fast trading, *Journal of Financial Economics* 116, 292–313.
- Bloomfield, Robert, Maureen O'Hara, and Gideon Saar, 2005, The “make or take” decision in an electronic market: Evidence on the evolution of liquidity, *Journal of Financial Economics* 75, 165–199.
- Brogaard, Jonathan, Allen Carrion, Thibaut Moyaert, Ryan Riordan, Andriy Shkilko, and Konstantin Solokov, 2018, High frequency trading and extreme price movements, *Journal of Financial Economics* 128, 253–265.
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2014, High frequency trading and price discovery, *Review of Financial Studies* 27, 2267–2306.
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2018, High frequency trading and the 2008 short-sale ban, *Journal of Financial Economics* 124, 22–42.
- Budish, Eric, Peter Cramton, and John Shim, 2015, The high-frequency trading arms race: Frequent batch auctions as a market design response, *Quarterly Journal of Economics* 130, 1547–1621.
- Cao, Charles, Oliver Hansch, and Xiaoxin Wang, 2009, The information content of an open limit order book, *Journal of Futures Markets* 29, 16–41.
- Carrion, Allen, 2013, Very fast money: High-frequency trading on the NASDAQ, *Journal of Financial Markets* 16, 680–711.
- Chaboud, Alain, Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega, 2014, Rise of the machines: Algorithmic trading in the foreign exchange market, *Journal of Finance* 69, 2045–2084.
- Chordia, Tarun, T. Clifton Green, and Badrinath Kottimukkalur, 2018, Rent seeking by low-latency traders: Evidence from trading on macroeconomic announcements, *Review of Financial Studies* 31, 4650–4687.
- Collin-Dufresne, Pierre, and Vyacheslav Fos, 2015, Do prices reveal the presence of informed trading? *Journal of Finance* 70, 1555–1582.
- Conrad, Jennifer, Sunil Wahal, and Jin Xiang, 2015, High-frequency quoting, trading, and the efficiency of prices, *Journal of Financial Economics* 116, 271–291.
- Cont, Rama, Arseniy Kukanov, and Sasha Stoikov, 2014, The price impact of order book events, *Journal of Financial Econometrics* 12, 47–88.
- Danielsson, Jon, Hyun Shin, and Jean-Pierre Zigrand, 2012, Endogenous extreme events and the dual role of prices, *Annual Review of Economics* 4, 111–129.
- Engelberg, Joseph, Adam Reed, and Matthew Ringgenberg, 2012, How are shorts informed?: Short sellers, news, and information processing, *Journal of Financial Economics* 105, 260–278.
- Fleming, Michael, Bruce Mizrahy, and Giang Nguyen, 2017, The microstructure of a U.S. Treasury ECN: The BrokerTec platform, *Journal of Financial Markets* 40, 2–22.
- Foucault, Thierry, 1999, Order flow composition and trading costs in a dynamic limit order market, *Journal of Financial Markets* 2, 99–134.
- Foucault, Thierry, Johan Hombert, and Ioanid Rosu, 2015, News trading and speed, *Journal of Finance* 71, 335–382.
- Foucault, Thierry, Roman Kozhan, and Wing Wah Tham, 2017, Toxic arbitrage, *Review of Financial Studies* 30, 1053–1094.

- Glosten, Lawrence, and Paul Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.
- Goettler, Ronald, Christine Parlour, and Uday Rajan, 2009, Informed traders and limit order markets, *Journal of Financial Economics* 93, 67–87.
- Hasbrouck, Joel, 1991a, Measuring the information content of stock trades, *Journal of Finance* 46, 179–207.
- Hasbrouck, Joel, 1991b, The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies* 4, 571–595.
- Hasbrouck, Joel, 1995, One security, many markets: Determining the contributions to price discovery, *Journal of Finance* 50, 1175–1199.
- Hautsch, Nikolaus, and Ruihong Huang, 2012, The market impact of a limit order, *Journal of Economic Dynamics and Control* 36, 501–522.
- Hirschey, Nicholas, 2017, Do high-frequency traders anticipate buying and selling pressure? Working paper, London Business School.
- Hoffmann, Peter, 2014, A dynamic limit order market with fast and slow traders, *Journal of Financial Economics* 113, 156–169.
- Huang, Roger, 2002, The quality of ECN and market maker quotes, *Journal of Finance* 57, 1285–1319.
- Jovanovic, Boyan, and Albert Menkveld, 2015, Middlemen in limit-order markets, Working paper, New York University.
- Kaniel, Ron, and Hong Liu, 2006, So what orders do informed traders use? *Journal of Business* 79, 1867–1913.
- Kirilenko, Andrei, Albert Kyle, Mehrdad Samadi, and Tugkan Tuzun, 2017, The flash crash: High-frequency trading in an electronic market, *Journal of Finance* 72, 967–998.
- Korajczyk, Robert, and Dermot Murphy, 2019, High frequency market making to large institutional trades, *Review of Financial Studies* 32, 1034–1067.
- Kyle, Albert, 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.
- Madhavan, Ananth, 1992, Trading mechanisms in securities markets, *Journal of Finance* 47, 607–641.
- Menkveld, Albert, 2013, High frequency trading and the new market makers, *Journal of Financial Markets* 16, 712–740.
- Rosu, Ioanid, 2019, Fast and slow informed trading, AFA 2013 San Diego Meetings; Paris December 2015 Finance Meeting EUROFIDAI - AFFI; HEC Paris Research Paper No. FIN-2015-1123.
- Thompson, Samuel, 2011, Simple formulas for standard errors that cluster by both firm and time, *Journal of Financial Economics* 99, 1–10.
- van Kervel, Vincent, and Albert Menkveld, 2019, High-frequency trading around large institutional orders, *Journal of Finance* (forthcoming).
- Zhu, Haoxiang, 2013, Do dark pools harm price discovery? *Review of Financial Studies* 27, 747–789.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication Code.